ALMA Correlator Preliminary Design Review

February 17, 2000
John C. Webber

Introduction

On January 20-21, 2000, the ALMA Correlator Preliminary Design Review was held in Charlottesville, VA in the NRAO auditorium. The purpose of the review was to examine plans for development and construction of a digital correlator for ALMA, including the design which was begun by the US MMA project and plans for a possible future European correlator.

Attendees for at least some part of the review were:

NRAO-CV:	John Webber, Ray Escoffier, Chuck Broadwell, Joe Greenberg, Bob Treacy, Jim Pisano, Bob Brown, Al Wootten, Paul Vanden Bout, Fred Schwab, Skip Thacker, Marc Rafal
NRAO-Tucson:	Darrel Emerson, Larry D'Addario
NRAO-Socorro:	Dick Sramek, Dan Edmans, Jim Jackson
OVRO:	Dave Woody
Haystack:	Alan Whitney
Europe:	Alain Baudry, Marc Torres, Albert Bos
Japan:	Yoshihiro Chikada
Canada:	Peter Dewdney, Brent Carlson
NSF:	Bob Dickman

Minutes of January 20: review of NRAO development

Introduction

Webber introduced the review by discussing the PDR objectives which were adopted by the US MMA project and are posted on the ALMA website (http://www.alma.nrao.edu/). There was mention of the fact that the ALMA project has not adopted official review guidelines, and that this should be addressed (as should a number of other organizational matters in the joint project).

Science Requirements

Wootten discussed the scientific goals and resulting requirements for sensitivity, bandwidth, frequency resolution, and time resolution. There was some discussion of these points. Woody wanted to know how much resolution would really be useful. Dewdney mentioned the dynamic range requirement, which may be linked to calibration accuracy. Wootten stated that the specification on this was 40 dB. Woody inquired if the proposed 16 msec dump time was for just one correlator phase; Emerson replied that the NRAO plan is to use offset front end frequencies rather than phase switching with Walsh functions due to the large number of dumps required for a complete cycle with 64 antennas. D'Addario noted that suppression of common-mode and other effects would be accomplished by 180-degree phase switching at the antenna, removed at the output of the digitizer. There was some discussion of whether this would be adequate. Torres remarked that it would be useful for someone to try this with an existing array. Chikada-san stated that NRO is planning to do this, with results expected in about 2 years. After some additional discussion, it was concluded that this was primarily a systems issue with little direct impact on the correlator, and the question of phase switching was relegated to off-line discussion.

NRAO Correlator Architecture

Escoffier then presented an overview of the NRAO correlator architecture. He explained one as-yet undocumented feature: signal distribution at the signal receivers in each rack will be via a replaceable paddleboard on the rear of each motherboard. Paddleboards for the first quadrant of the correlator will be configured so that all IF signals from up to 32 antennas can be processed. When all quadrants are installed, replacement paddleboards will enable full operation with all 64 antennas. There was some discussion of how to generate the analog sum for phased array mode. It was agreed that even the largest bandwidth could be assembled with sufficient logic and memory, but D'Addario pointed out that no seriously proposed VLBI recording system approaches the sort of bandwidth of the ALMA IF system. Albert Bos noted that, at Westerbork, there are two tied array adder outputs which can be used for VLBI or for high spectral resolution tied array interferometry.

Digitizer Requirements

D'Addario presented the present understanding of requirements for the digitizer. Dewdney asked for the justification for the specifications on threshold voltages. D'Addario replied that it was mostly from practical arguments about how accurately one may expect the thresholds to be set; the suggested specification should result in an SNR loss of 1-2%, but this has not been simulated. Webber noted that, at the recent URSI meeting, Don Backer raised the question of whether the non-linear quantization correction has been worked out for the sample-filter-requantize mode; it has not, but in principle can be calculated or at least simulated. Albert Bos questioned whether the fine delay could really be common for all digitizers at a given antenna, because of effects on phased array and calibration modes. Dewdney noted that the fine delay is tricky and at the VLA is not quite right, but that it could be expensive to do each digitizer separately. Bos noted that one bit of bulk delay in the digitizer would reduce the loss; Escoffier and D'Addario replied that it would gain only a factor of 2.

European Digitizer Plans

Baudry discussed European plans for digitizer development, which the US correlator group understood to be an exclusively European responsibility. Baudry noted that, if two serial quantization losses really result in multiplication of the losses from each stage, that there would be a 40% increase in observing time required vs. an analog system, the exact value depending on the number of levels in the quantizations [this was discussed again later]. The European digitizer group has been formed and, as a first step, has looked into commercially available A/D converters, of which one goes as high as 4 Gsamp/sec and several achieve 1 to 2 Gsamp/sec. The group is looking for engineering contacts within one the companies which make these devices. The European group has made a survey of fast digitizers fabricated or planned in various radio observatories. Dewdney noted that the AT intends to use an InP-based digitizer. D'Addario noted the important fact that none of the commercial devices has an analog response which goes to 4 GHz, so they could not be used with the planned 2-4 GHz bandpass of the ALMA IF system, although a change in the IF system to 0-2 GHz could work. There was some discussion of the circuit details of the digitizer module. Woody proposed that demultiplexing on the ASIC could be used, but Torres stated that this would probably result in excessive digital feedback into the analog part of the chip, with undesirable results. Carlson wanted to know if limiting the clock slew rate would help the feedback problem. There was some additional discussion of how to achieve isolation. Dewdney noted that commercial chips all demultiplex by 2X or more on-chip; Torres noted that none of them does bandpass sampling.

D'Addario stated it was easier to bandpass sample if the sample-and-hold comes before the analog-to-digital stage; Baudry said he hadn't thought about it. Webber again raised the issue about the effects on SNR of digitization, filtering, and re-quantization. D'Addario noted that it can be calculated in principle analytically, but at least by Monte Carlo simulation, and that the effect would have to be determined only once. He contended that 2-bit followed by 2-bit was not as bad as two serial losses would naively imply [see further discussion on January 21]. Baudry had stated that the European group's plans were first to get an operational unit which would satisfy the specifications with 2-bit, 4-level quantization before proceeding to the 3-bit case; D'Addario wanted to know the justification; Baudry replied that he thought the 3-bit operational case would be very difficult. Webber asked when prototypes would be available for use with the correlator. Baudry replied that the European group had not actually discussed the question of prototype modules, but only the development of the critical digitization submodule. He was not prepared to commit to definite plan or schedule, because the group had not addressed these issues. He stated that the European group thought they needed a usable pre-prototype before even thinking about other aspects such as system-related or packaging issues. It became clear that there had been some misunderstanding between the US and European groups, in that the US had assumed that the Europeans had taken full responsibility for timely digitizer module development, but the Europeans had assumed they were undertaking a long-term pre-prototype effort. Webber stated that it was highly desirable to get a digitizer at the time that a prototype correlator became available; Woody contended that the correlator could be tested adequately using only PN sequence generators. Sramek noted that the schedule calls for a commitment to quantity production in April 2003, which according to the presented European development schedule could not be met for the digitizer. Dewdney thought that several milestones should be generated for the digitizer development. There was some discussion of what end-to-end testing of the prototype interferometer really means and what elements should be present during that

testing. In response to a question from Dewdney, D'Addario discussed why bandwidth sampling would be better--it makes the IF system easier to build with good band edges. It was clear at this point that there was a conflict in specifications, because the European group has been assuming that baseband sampling would be used. It was agreed that this should be further discussed at the system level.

Signal Interfaces

Escoffier discussed current NRAO plans to use 125 MHz signal interfaces: for chip-to-chip, unterminated; for card-to-card, terminated using Xilinx Virtex-E drivers; for rack-to-rack, LVDS on shielded Twin-Ax cables. Additional options include operating the interfaces at 250 MHz or, with a change to Altera parts, even 500 MHz. Either option would reduce the number of cables but possibly make it more difficult to produce reliable communications. Woody suggested using commercially available unshielded twisted pairs, which are used in RAID arrays and are about 1/3 the price of Twin-Ax. After some discussion, the fact that signals branch out to many places instead of having many signals go to the same place appeared to make it impossible to use the inexpensive mass-produced cable assemblies. Escoffier discussed the use of quadrature capture in the signal interfaces, replying to a Torres question. Each interface would be tested, and the error-free phase would be captured by software for reconfiguration at power-up. Woody wanted to know if all of these would be individually tested; Escoffier thought it would certainly be done for all rack-to-rack interfaces, but it might not be needed for all card-to-card interfaces. The goal is zero errors on all interfaces during test sequences.

FIR Filter

Escoffier discussed the NRAO FIR digital filter. Carlson inquired if it would be possible to use a symmetric 127-tap filter instead of 128 taps; Escoffier replied that this could not be accommodated in the FPGA implementation. Woody asked if it would not be preferable to use a Verilog/VHDL description instead of schematic capture for the FPGA personality. Escoffier replied that the NRAO engineers are very familiar with the use of Xilinx schematic capture tools; Carlson noted that those tools allow the schematics to be ported to new Xilinx chips. There was some discussion of multiple-spectral-line restrictions; essentially, it is possible with the NRAO design to do each of 4 pairs independently, but that is the maximum number of passbands available.

Station Card

Escoffier presented the architecture and functions of the station card. In response to questions from Dewdney, he noted that all monitoring is performed through the microprocessor port and that it has not been determined how often the state counters will be read. Torres asked about blanking efficiency, and Escoffier noted that due to newly-adopted double-buffered memory scheme, the number of bits lost from each time slice of 62500 would become 256, or a net loss of 0.4%.

Fiber Optics

Jackson discussed the plans for fiber optic interfaces and digital data transmission. There was considerable discussion of setting and tuning the wavelengths of multiple lasers. Woody wanted to know why wavelength-division multiplexing was being considered instead of using multiple fibers. Edmans noted that it made keeping delays uniform a lot easier. D'Addario noted that there was an absolute need to make all the bits from one digitizer properly aligned, although they will be separated during transmission. There is also a question of cost. Sramek noted that minimizing the number of fibers also minimizes the size of the patch panel needed to support multiple antenna foundations. The question of how much power could be used per fiber without causing non-linear problems, and Edmans replied that 1 mW per laser with the planned number of wavelengths per fiber would be OK. Woody stated that he was still worried about the effect of laser frequency drift on delay. Dewdney wanted to know if header and timing information would be added. Jackson stated that he was studying the alternatives, such as making everything self-synchronizing. Baudry inquired about crosstalk effects. Edmans replied that simulations show there should be no significant optical problem with this design, at least up to 16 optical channels per fiber. Torres wanted to know how to handle different phase on different signals; Jackson indicated that buffering in FPGAs would synchronize all outputs to the system clock, using a barrel shift scheme. There was some discussion about buffering and timing, which left unresolved some disagreements about clock recovery and synchronization among different signals.

Correlator Card and Custom Chip

Greenberg discussed the architecture of the correlator card and the custom chip, along with various alternatives in chip development. D'Addario wanted to know about the limit on power consumption per rack. Escoffier stated that he had no absolute number in mind, and some discussion of practicality ensued. Carlson noted that, for the JCMT correlator, heat spreaders were needed to minimize temperature gradients. Whitney stated that the SMA correlator consumed 3W per chip and 2500W/rack, for which they have installed an overengineered cooling system. Rafal noted that 0.18 micron chips may not tolerate as high a gate temperature as 0.25 micron chips. D'Addario asked about the best way to exploit the 0.18 micron technology, noting that several alternatives were possible: reduce power, increase number of lags, accommodate more antennas per chip, or increase the clock rate. Escoffier stated that the frequency limit for this level of power dissipation is believed by the chip designer to be ~150 MHz. Changing the matrix size might require re-bidding the chip design. The voltage has already been lowered to support ~150 MHz maximum speed. Jackson advised exhaustive and careful characterization of any chip designed to run at a different voltage from that commonly used for a given process. There was a discussion of what constituted adequate levels of characterization, simulation, and application of relevant experience with different processes. Webber raised the question of whether doubling the frequency resolution by going to an 8K chip would be worthwhile scientifically. Woody and Wootten concluded that the major effect would be to increase the speed at which wideband searches for narrow extragalactic spectral lines could be conducted, and Woody stated that about half the time on the OVRO array would used in the widest band mode when their new correlator is implemented. Dewdney returned to details about the chip by asking if it would be useful to carry out a formal analysis of the probability of success of various chip performance alternatives, and in doing this to include batch-to-batch variation. Whitney stated that it is nearly impossible to get information from any foundry about

this factor; it is considered highly proprietary. He also noted that it was essential to keep track of chips on a wafer-by-wafer basis. There was discussion of details of heat dissipation and other chip design factors. Torres noted that the blanking efficiency would further decrease with an 8K chip. D'Addario noted that there would also usually be some loss of bits associated with geometry and phase switching.

LTA and LTA Adder Tree

Broadwell discussed the architecture and function of the long-term accumulator and adder tree. In response to several questions, he clarified several points about transfer mode, speeds, and the number of spectral points in various modes. Woody noted that the resolution with a full 16 GHz bandwidth and the 4K chip is 15 MHz. This is close to the ideal of 10 MHz for highly-redshifted CO searches, and would be adequate for the task, although double this resolution would be useful (although it would double the output data rate). Rafal wanted to know if a weighted analysis of various observing programs would be possible or useful; there was some discussion about the uncertainties of future needs, and there was no definite conclusion reached. D'Addario wanted to know if longer integrations in autocorrelation mode were needed; some discussion resulted in the conclusion that the present design was OK. However, the subject of binning results as opposed to merely buffering sequential dumps was raised and discussed at some length. It was concluded that, although the capability of binning data before output is not absolutely required, it is a highly desirable function to include at design time if it can be accommodated without significant schedule or cost impact. The NRAO correlator group agreed to study this question seriously. Escoffier noted that more control parameters would be required, and that there could be implications for subarraying.

Computer Interface

Broadwell presented the means by which the accumulated data would be transferred to the data computer system. There were a number of questions which were finally answered by clarifying the plan to allow up to 1024 individual block readout requests to be transferred to the LTA in each request message. Woody noted that, for pointing calibration, only a few lags are required, but they are needed at a high rate. Broadwell replied that providing anything except the full 512 lags (the current plan) would require a modification of the request word and a change in the logic which serves requests, and that this should be easy to do in the hardware. There was some discussion of the choice of microprocessors, which became controversial and was relegated to off-line discussions.

Control Computer

Pisano discussed the computer system which will receive and process data from the LTA. Woody asked for clarification of several issues related to the use of the CAN bus, and there was a general discussion which resulted in the conclusion that the CAN bus is probably OK for the correlator. The strawperson design of a post-processor using Sky equipment was discussed at length. Issues raised were the use of DSPs instead of general-purpose machines and the reorganization of data in software. One conclusion reached was that, although the leads and lags from the same quadrant of the correlator could conceivably be made contiguous in hardware, this

would be difficult for those modes in which different quadrants process different lags of the same IF signals.

Minutes of January 21: review of NRAO development, concluded

Quantization Noise

D'Addario presented arguments concerning the issue of sampling-filtering-requantizing. He noted that, in the direct pass-through of 8-level to 4-level data (or even 4-level to 4-level), there is no additional loss whatsoever. He argued that filtering rejects first stage quantization noise to a large degree, at least in the case of low-pass filtering. Bos argued that there would still be quantization noise inside the retained passband, and that just rejecting the aliased noise was insufficient to produce net loss of SNR. Rafal noted that this should all be amenable to simulation; D'Addario thought it would be tricky and lengthy. There was some discussion of how to interpret simulations, and general agreement that someone should do it.

Discussion of European correlator plans and other points

Baudry summarized European thinking about a future, larger correlator. The European correlator group has been thinking that there would exist a branching point in 2004; they now perceive a conflict due to an earlier start and good progress in the NRAO effort. They would like to see a merger of efforts, possibly by forming a joint design group. There are many interested organizations in Europe; several of them are getting close to signing a formal organizing agreement. The present European plan is for a one-year feasibility study throughout 2000, with further work in 2001 resulting in a preliminary design report at the end of 2001. Various groups are looking at possibilities for FIR filters and digital image rejection mixers. Webber asked why a future correlator was needed at all, given that the NRAO design appears to satisfy all the science requirements for ALMA. It was at this point it became clear that there was insufficient communication across the Atlantic: the Europeans have been assuming that the NRAO design would accommodate only the first 36 antennas (the plan of the US Reference Project) and that a new design would be needed for the final array of 64 antennas. There was some discussion of this point. D'Addario asked when we needed to decide whether to build the full 64-antenna NRAO correlator or just one quadrant (OK until the 33rd antenna arrives in Chile), and stated it would be desirable to make the decision now. Simon inquired whether there might be some science which the present correlator would not be able to do; Wootten stated that the only thing he could think of was bistatic radar, which might be better served by a special purpose data processor. Chikada-san noted that the Japanese contribution may result in an array with more antennas, and that they are interested in studying a correlator to serve an even larger array than 64 antennas--perhaps up to 96 antennas. There was a discussion of possible future technology improvements and presently unknown science goals. Chikada-san noted that all partners in ALMA want to develop things and get experience with new technology. Webber replied that, although this was indeed an important consideration, it was also necessary to follow a plan which assures that the correct correlator is available for use when antennas begin arriving in Chile; this is satisfied by the present NRAO development effort.

There was a discussion of image generation and resulting data compaction, which was relegated to off-line discussions.

Returning to the correlator development schedule, there was a discussion of what decisions need to be made, and when they must be made. There was a proposal to call the first quadrant of the NRAO correlator the "32-antenna" correlator, since it will be configured to handle all IF pairs for up to 32 antennas. Sramek questioned the incremental cost of a split chip procurement, which would have the effect of further postponing the date at which a final decision would need to be made regarding building the 64-antenna correlator. It was concluded that this could increase chip procurement cost by a factor of 1.5 to 2, which is significant. D'Addario wanted to know what constitutes a greater form of waste, and argued that all efforts except one should be abandoned, and the money spent on something else useful to ALMA. There was some discussion of pricing and schedule. Brown asked if we ought not to set a breakpoint for the decision regarding how many quadrants to produce. Rafal thought the real question is whether anything concrete will happen in Europe or Japan in the next two years, before we need to decide on the size of the chip procurement. Sramek reinforced D'Addario's point that we need to know the criteria on which that decision will be based. D'Addario argued that we should decide the ultimate size of the NRAO correlator right now, and start a new correlator much later; aiming at a useful life of 10 years for this correlator would not be unreasonable. Chikada-san noted that a few years ago the different partners were thinking of separate arrays which would occasionally be combined for some periods. However, we are now looking at a single array and so need to look at developments in the US, Europe, and Japan which will result in a long-term science benefit. Woody asked if the ALMA operating budget would support continuing development of upgrades such as new correlators. This was thought to be highly speculative, and advantages were seen in having at least the first correlator upgrade perceived as part of the construction effort. Dewdney stated that, in his opinion, there was no doubt that the NRAO correlator as planned will satisfy the stated science requirements. However, he expects that studies of different architectures will show ways to achieve more capacity per dollar; he therefore argued that the decision point on the size of the NRAO correlator build should be in 2002. D'Addario noted with chagrin that the participants seemed to be implying that the NRAO correlator is too early, and that in his opinion this criticism is unjustified and should not be taken to imply that we should wait. Dewdney, Woody, and others replied that this was not the case at all, that the time scale looks good, and that development must continue at the present rate or we will have nothing in 2004 when the first antennas arrive in Chile. Whitney stated that there had been excellent progress on the design, and that it must continue. Woody noted that the NRAO design is very cost effective, but it has the feature that it cannot be easily expanded and we have to keep that in mind as the project changes.

Open Discussion

Dewdney asked if the question of re-sorting the data before sending it to the computer system could be re-examined.

There was some discussion of chip procurement. Escoffier noted that there were additional costs beyond the chip itself, such as power distribution.

There was a discussion of the FIR filter location and the question of 3- or 4-bit digitization. Jackson guessed that sending 3 bits per sample rather than 2 would result in an increased cost for the fiber optic system of roughly $3.6M at today's prices. Bos noted that this could be worth the investment, because bringing 3 bits per sample to the correlator gives future flexibility. Escoffier wanted to see some firmer cost estimates for the various possible configurations before a decision is made. D'Addario would like to study the SNR implications as well. He also noted that since a commercial standard of 10 Gbit/sec is becoming established, and we are using only 8 Gbit/sec per link, we are throwing away 20% of what we pay for. Carlson stated that the data will need lots of transitions, and one way to do this is to adopt 8/10B encoding; others stated that simply multiplying the raw data with a PN sequence would be adequate. Escoffier asked again if we could set a date for the FIR filter decision; there was no conclusion, but this was passed off to the systems group for consideration.

Committee session and recommendations

The official committee agreed that they were impressed with progress on the NRAO design and the thoroughness with which it has been carried out.

With regard to the US MMA project guideline questions, the following discussions and comments resulted:

1. Are the top level performance requirements for the subsystem complete and adequate?

Yes, the requirements have been specified with the possible exception of a few items which should be examined. The question of binning should be looked at carefully. D'Addario noted that the correlator does not support sideband separation, which is accomplished by frequency offsets. Woody asked that the details of the correlator and its capabilities should be distributed to experienced interferometrists for comment.

2. Have the correct design solutions been selected for study and development during the MMA D&D Phase? Are there important alternate solutions that are not being studied?

It was agreed that the NRAO design is a good solution for the correlator requirement. The one area identified as needing increased attention is the digitizer development, which was discussed at length. Whitney and Baudry stated that this was a risky area which needs a lot of work, and that success on the required time scale is not guaranteed. Sramek asked if more resources needed to be put into the digitizer effort on one or both sides of the Atlantic. Baudry stated that he had thought the US had an alternate plan for the digitizers which was being pursued. Emerson and D'Addario both said that active US development had been put on hold because they mistakenly believed the European side had agreed to do the digitizer, but starting up that effort again could be considered. Woody noted that bandpass sampling increases both cost and risk, and that a commercial product might be found which would work for baseband sampling, albeit with more bits than we need and much higher power than a specialized solution. Whitney

urged that we examine the consequences of such an alternate solution. D'Addario noted that we might consider going to more digitized channels, each with smaller bandwidth. He would be willing to work on the project, given technician time and money. It was agreed that any restart of an NRAO effort should wait until the European group has had time to re-examine their plan in the light of new information; Baudry will present a report at the time of the February 28 systems PDR. Finally, Rafal noted that the requirements still needed to be refined and urged that some effort be put into this.

3. If a major procurement is required for the subsystem during the MMA D&D Phase, is the procurement plan correct?

The only area identified as a major procurement item was the custom correlator chip. There was a unanimous recommendation that the change from 0.25 to 0.18 micron technology be carried out, provided that the ongoing study of consequences resulted in confirming the conclusions which are now believed correct: lower power and lower overall cost. It was recommended that the decision to go to an 8K chip not be made until a detailed evaluation of the power dissipation of the 4K, 0.18 micron chip was known.

Whitney stated that there should be a well-defined decision point for the full correlator construction decision; the Europeans have essentially the same amount of money as the US and a strong interest in conceptual designs; correlators present a significant intellectual challenge in the ALMA project; a European contribution is needed; therefore, the Europeans should have the opportunity to propose a new design before a commitment is made for the full 64-antenna US correlator. Rafal and Woody both agreed that it was OK to wait until 2002-2003 to make a final decision which includes all factors.